Swarm AI to Detect Deceit in Facial Expressions

Louis Rosenberg and Chris Hornbostel
Unanimous A.I.
San Francisco, CA, USA
Louis@Unanimous.ai

Niccolo Pescetelli
University of Oxford
Clarendon, UK
niccolo.pescetelli@chch.ox.ac.uk

***Abstract*—In the natural world, many species amplify their intellectual abilities by working together in closed-loop systems. Known as Swarm Intelligence (SI), this process has been deeply studied in schools of fish, flocks of birds, and swarms of bees. The present research employs artificial swarming algorithms to create "human swarms" of online users and explores if swarming can amplify the group's ability to detect deceit. Researchers recruited 168 participants and divided them randomly into five online swarms, each comprised of 30 to 35 members. Working alone and in networked groups, participants were given tasks with evaluating a set of 20 video clips of smiling people. Each video clip depicted either 1) an *authentic smile* generated in response to a humorous cue; or 2) a *deceitful smile* generated falsely upon command. Across the population of 168 participants, the average individual incorrectly identified the deceitful smiles in 33% of the trials. When making evaluations as real-time swarms, the error rate dropped to 18% of trials. This large reduction in error rate suggests that by swarming, human groups can significantly amplify their ability to detect deceit in facial expressions. These results also suggest that swarming should be explored for use in amplifying other forms of social intelligence.**

Keywords—Swarm intelligence; artificial swarm intelligence; collective intelligence; human swarming; artificial intelligence

I. INTRODUCTION

Across the natural world, many social species have evolved methods for amplifying their collective intelligence by forming real-time closed-loop systems. From flocks of birds and schools of fish, to swarms of bees and colonies of ants, this amplification of intelligence has been observed across a wide range of natural systems. Biologists generally refer to the phenomenon as Swarm Intelligence (SI). In recent years, the principal of Swarm Intelligence has inspired a new category of A.I. research that uses swarming algorithms to form real-time closed-systems among online human groups with the objective of amplifying intelligence beyond the abilities of the individual members. Known as Artificial Swarm Intelligence (ASI), the technique enables distributed human groups to combine their knowledge, wisdom, and intuitions into unified system that can answer questions, make predictions, and solve problems by converging on solutions in synchrony [1], [2].

Recent studies show that human swarming can significantly amplify the predictive ability of online groups, outperforming individual members. In one recent study conducted by researchers at Unanimous A.I. and Oxford University, human subjects were tasked with predicting a set of 20 official Las Vegas wagers known as "proposition bets" on Super Bowl 50. Traditional crowd-sourced predictions were pitted against the predictions made by real-time swarms. The crowd was composed of 467 football fans who provided their predictions by online survey. The swarm was composed of 29 football fans who provided their predictions together as a closed-loop system. Although the crowd was 16 times larger than the swarm, it was far less accurate, achieving only 47% correct predictions and generating a 9% gambling loss. The swarm achieved a significant improvement, producing 68% correct predictions and generated a 36% gambling win [3].

While many recent studies have demonstrated the ability of human swarms to amplify the predictive intelligence of human groups, no prior work has looked at the social intelligence of real-time swarms. To address this, the present study tested the ability of human swarms to identify deceit in human faces. Specifically, the study compared the ability of individuals with the ability of swarms when assessing if a person in a video clip was producing an "authentic smile" (i.e. a genuine smile evoked in response to a joyful stimuli), or a "deceitful smile" (i.e. a forced smile evoked on demand).

II. JUDGING HUMAN DECEIT

In 1969, Ekman and Friesen published the first critical research linking facial cues and human deception. They defined "leakage cues" as involuntary expressions that reveal if an individual's true feelings don't match what they're consciously attempting to convey [4]. Among common facial expressions, smiles have been identified as a significant leakage cue that can be used to determine if a person is being honest or deceitful, especially if they are deliberately faking their positive emotions [5].

While the difference between genuine smiles and deceitful smiles can be established by trained researchers performing rigorous analysis of smile features, most people are not very good at telling the difference during live experiences. That's because the facial cues are often subtle and easily missed by the untrained eye [6]. In a recent study, a group of 217 subjects were asked to view a set of 20 videos of individuals smiling. The videos represented a mix of genuine smiles (produced by enjoyment) and fake smiles (produced on demand). Across the 217 subjects, the average person's assessment was incorrect 32% of the time [7].

The question thus remains: can artificial swarming be used to improve the ability of human groups to detect deceit when evaluating facial expressions? If so, human swarming may be an effective technology for amplifying many forms of social intelligence.

978-1-5386-2823-2/17/$31.00 ©2017 IEEE

III. Swarms as Intelligent Systems

Research into human swarms has been inspired by the intelligence of natural swarms, which serve as the basis for most swarming algorithms. The present research was modeled after the decision-making in honeybee swarms, as it's been observed to be remarkably similar to decision-making in neurological brains [8], [9]. Both employ large populations of simple excitable units (i.e., bees and neurons) that work in parallel to integrate noisy evidence, weigh competing alternatives, and converge on decisions in synchrony. In both, outcomes are arrived at through a real-time competition among sub-populations of excitable units. When one sub-population exceeds a threshold level of support, the corresponding alternative is chosen. In honeybees, this enables optimal decisions over 80% of the time [10]-[12]. It is this amplification of intelligence that Artificial Swarm Intelligence aims to enable among distributed networked humans.

The similarity between neurological intelligence and swarm intelligence becomes even more apparent when comparing decision-making models that represent each. For example, the decision-making process in primate brains is often modeled as mutually inhibitory leaky integrators that aggregate incoming evidence from competing neural populations. A common framework is the Usher-McClelland model [13] represented in Fig. 1 below. This can be directly compared to swarm-based decision models, like the honey-bee model in Fig. 2 below. As shown, these swarm-based decisions follow a similar process, aggregating incoming evidence from sub-populations of swarm members through mutual excitation and inhibition.

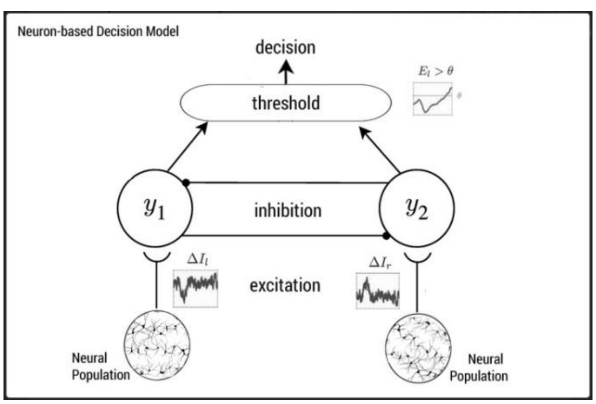

Fig. 1. Usher-McClelland model of neurological decision-making.

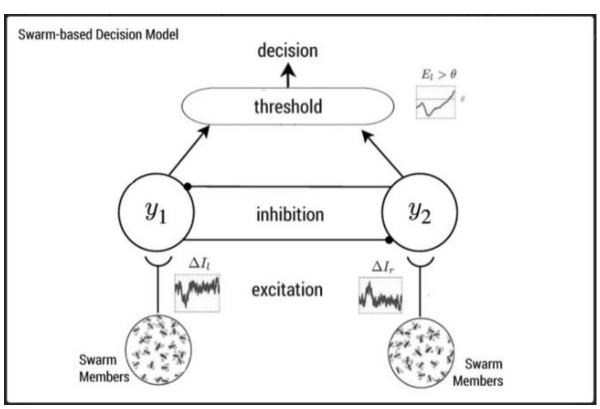

Fig. 2. Mutually inhibitory decision-making model in bee swarms.

IV. Enabling "Human Swarms"

Unlike many other social species, humans have not evolved the natural ability to form a closed-loop Swarm Intelligence. That's because we lack the subtle connections that other organisms use to establish tight-knit feedback-loops among members. Schooling fish detect vibrations in the water around them. Flocking birds detect motions propagating through the group. Swarming bees use complex body vibrations called a "Waggle Dance". Thus to enable a real-time Artificial Swarm Intelligence among groups of networked humans, specialized technology is required to close the loop among members.

To address this need, an online platform called **UNU** was developed by Unanimous A.I. in 2015 to allow distributed groups of users to login from anywhere around the world and participate in a closed loop swarming process [14], [15]. Modeled after the closed-loop decision-making of honeybee swarms, the "Swarm A.I." algorithms employed by the UNU platform allows groups of independent actors to work in parallel to 1) integrate noisy evidence; 2) weigh competing alternatives; and 3) converge on final decisions in synchrony, while also allowing all participants to perceive and react to the changing system in real-time, thereby closing a feedback loop around the full population of participants.

As shown in Fig. 3, participants in the UNU platform answer questions by collectively moving a graphical puck to select among a set of alternatives. Each participant provides input by manipulating a graphical magnet with a mouse or touchscreen. By positioning their magnet, users impart their personal intent on the puck. The input from each user is not a discrete vote, but a stream of vectors that varies freely over time. Because the full population of users can adjust their intent at every time-step (200 ms), the puck moves, not based on the input of any individual, but based on the dynamics of the full system. This enables real-time physical negotiation among all members, empowering the group to collectively explore the decision-space and converge on the most agreeable solution in synchrony.

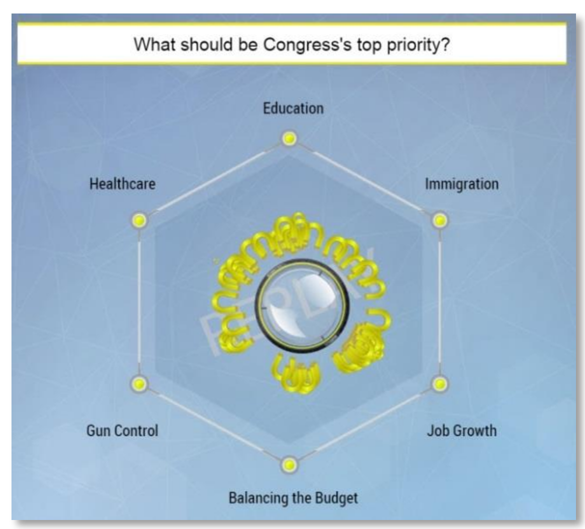

Fig. 3. A human swarm answering a question in real-time.

978-1-5386-2823-2/17/$31.00 ©2017 IEEE

It is important to note that participants do not simply vary the direction of their input, but also modulate the magnitude of their input by adjusting the distance between the magnet and the puck. Because the puck is in continuous motion across the decision-space, users need to continually move their magnet so that it stays close to the rim of the puck. This is significant, for it requires participants to be engaged continuously during the decision process, evaluating and re-evaluating their personal contribution. If they stop adjusting their magnet to the changing position of puck, the distance grows and their applied force wanes. Thus, like bees vibrating their bodies to express sentiment in a biological swarm or neurons firing activation signals to express sentiment in a neural-network, the participants in an artificial swarm must continuously express their changing preferences during the decision process, or lose their influence over the collective outcome.

V. DECEIT ASSESSMENT STUDY

To address whether an "human swarm" comprised of distributed online participants can more accurately assess the authenticity of smiles compared to individual human assessors, a formal research study was conducted. A pool of 168 human participants were recruited online, each paid approximately $2 for their time. The population was divided randomly into five online swarms, each comprised of 30 to 35 members. Working alone and in swarms (via UNU), the participants were tasks with evaluating a set of 20 video clips. Each clip was a headshot of a smiling person that was 3 seconds in duration and depicted either 1) an *authentic smile* generated in response to humorous cues; or 2) a *deceitful smile* generated falsely upon command. The videos were sourced from existing research by smile expert Paul Eckman [16].

Fig. 4 below shows snapshots from two of the smile videos. The snapshot on the left depicts a joyful smile that was generated in response to an authentic stimulus. The snapshot on the right depicts a fake smile that was produced deceitfully on demand, and not in response to a joyful stimulus.

Authentic "Joyful" Smile   Inauthentic "Fake" smile

Fig. 4. Shown are snapshot from two sample smile videos.

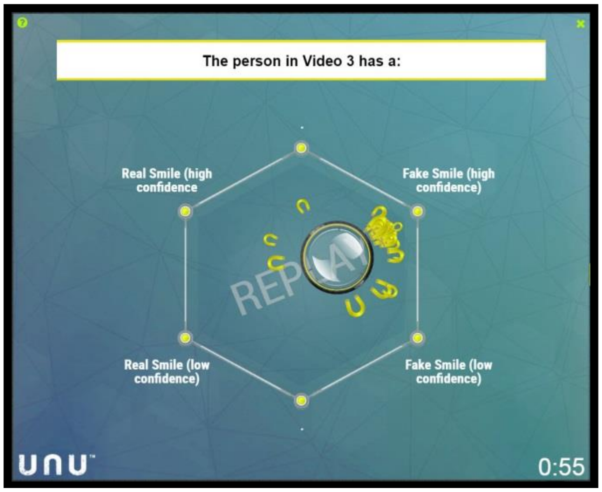

Fig. 5. Shown is a snapshot of a human swarm of 35 participants in the process of assessing a smile video in real-time.

Upon viewing each of the videos, each test subject was required to immediately express their personal assessment as to whether the video depicted a genuine "joyful smile" or a deceitful "fake smile." This testing performed twice for each individual participant – once by working alone and reporting their assessment on a standard online survey, and once by working as part of a real-time closed-loop swarm. When working as a swarm, the participants had no contact with other participants, as all were logged in from separate locations and were provided with no means of direct communication.

Fig. 5 above shows a snapshot of one of the five swarms in the process of assessing one of the 20 videos. As shown, all members of the 35 person group are working together to move the glass puck by individually positioning and repositioning their graphical magnets in synchrony. In this way, the group explores the decision-space and converges on a single unified assessment. It's important to note that although the image shows the full swarm of magnets in real-time, the individuals could only see their own magnet and the puck, but not the magnets of other members of the swarm.

It should be noted that each swarm was limited to only 60 seconds for viewing and assessing a single video, with most assessments being executed in under 20 seconds. It should also be noted that videos were only played once – subjects were not allowed to repeatedly review a video as they assessed the authenticity of the smile.

Data was collected for 168 individual participants, each providing a personal assessment of 20 videos. This produced a data set of 3360 smile evaluations performed by individual persons. Data was also collected for the five swarms, each

978-1-5386-2823-2/17/$31.00 ©2017 IEEE

providing a swarm assessment of 20 videos. This produced a data set of 100 swarm-based assessments for comparison against the 3360 individual assessments.

VI. RESULTS

As shown in Fig. 6, the raw data was processed for each of the five trials by computing the average number of errors for the participants in each trial and comparing the average error rate with swarm performance. Each trial was comprised of 20 video assessments, of which individuals averaged 13.4 correct smile assessments and 6.6 incorrect smile assessments. This corresponds with a **34%** error rate which conforms to prior research into human ability to assess deceitful smiles [7]. When those same individuals worked together as swarms, however, the average number of errors per trials was reduced to 3.6 incorrect assessments, which corresponds to an **18%** error rate. This is a significant improvement, corresponding with an average error reduction from swarming of **46%** across trials (±16%). In other words, by working together as real-time swarms, each group of participants was able to assess the authenticity of smiles with 46% fewer errors, on average, as compared to individuals working alone.

Although this study aims to compare the performance of human swarms to individual assessors, we can also compare the performance of the swarm to the traditional crowd-sourcing method of aggregating poll results across sets of independent respondents. Doing this for each of the five trials, the most popular poll result across the members of each group was used as the final smile-assessment for that group. This produced an average of **5.8** incorrect assessments across the five trials, which is still a significantly higher error rate than the 3.6 assessment errors for swarm-based responses. Specifically, swarming decreased the smile assessment errors by **34%** as compared to the traditional "Wisdom of Crowd" methodology.

To further assess statistical significance, we compared the swarm performance to the performance that would be expected by chance from a matching population using a bootstrap approach. For each twenty-video trial, we took a random sample of 20 individuals who participated in that trial, taking the first individual's assessment for the first video, the second individual's assessment for the second video and so on until we had 20 assessments from the 20 randomly selected individuals. We then averaged the accuracy of these samples. We repeated the procedure (i.e. random selection of 20 individuals and response assignment) 10000 times and computed the average distribution of correct answers for that trial.

	Group Size	INDIVIDUALS Avg # of Errors	SWARM # of Errors	Error Reduction
Trial 1	33	5.73	3	48%
Trial 2	35	7.03	4	43%
Trial 3	35	6.64	5	25%
Trial 4	30	6.40	2	69%
Trial 5	35	7.31	4	45%
Average	**34**	**6.62** σ = 0.61	**3.60** σ = 1.14	**46%** σ = 16%

Fig. 6. Results across five trials of the smile assessment test.

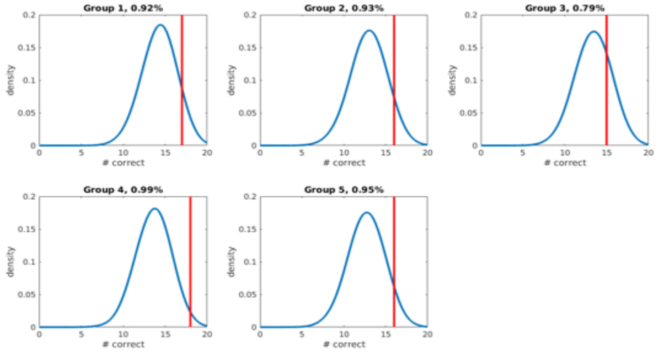

Fig. 7. Swarm outperforms individuals in Bootstrap analysis.

The density distributions, as shown in Fig. 7 by the curved blue lines, represent the average number of correct predictions that should be expected by chance in each trial by a matching assessor population. It can be seen that swarms, represented by the straight red lines, are well above the mean as compared to individual predictions in all trials. In fact the p-scores for all five trials were under 1%, indicating that the statistical chances of the human swarm outperforming the individuals in the smile assessment tasks was over a 99% probability.

VII. CONCLUSIONS

As expressed above, the results suggest that forming an Artificial Swarm Intelligence comprised of 30 to 35 online human participants can significantly reduce the error rate when assessing human smiles for deceit versus authenticity. This suggests that human swarms are not only useful for amplifying predictive intelligence, as shown by prior research, but are also useful for amplifying social intelligence, especially when social tasks in involve challenging subjective assessments. In addition, because smile authenticity is a direct indicator of human honesty, these results open a wide range of possible applications of ASI technology from intelligence screening to jury selection. Future research is required to further explore the ability of human swarms to identify deceit, not just in facial expression but in verbal cues and other social indicators.

ACKNOWLEDGMENT

Thanks to David Baltaxe for his efforts supporting the online data collection for this research study. In addition, thanks to Unanimous A.I. for the use of the online UNU platform for human swarming.

REFERENCES

[1] Rosenberg, Louis. "Human Swarms, a real-time method for collective intelligence." Proceedings of the European Conference on Artificial Life 2015, pp. 658-659

[2] Eberhart, Russell, Daniel Palmer, and Marc Kirschenbaum. "Beyond computational intelligence: blended intelligence." Swarm/Human Blended Intelligence Workshop (SHBI), 2015. IEEE, 2015.

[3] Rosenberg, Louis. Baltaxe, David and Pescetelli, Nicollo. "Crowds vs Swarms, a Comparison of Intelligence," 2016 Swarm/Human Blended Intelligence (SHBI), Cleveland, OH, 2016, pp. 1-4. doi: 10.1109/SHBI.2016.7780278

[4] Ekman, Paul. and Friesen, Wallace. The repertoire of non-verbal behavior: categories, origin, usage, and coding. Semiotica, 1. 49-98. 1969

[5] Ekman, Paul., Friesen, Wallace., and Davidson, Richard. The Duchenne Smile: Emotional Expression and Brain Physiology. Journal of Personality and Social Psychology, 1990, Vol 58, No 2, 342-353.

[6] Ekman, Paul., Friesen, Wallace., and O'Sullivan, Maureen. Smiles when lying. Journal of Personality and Social Psychology, 1988, Vol 54, 414-420.

[7] Kajdasz, James. A demonstration of the Benefits of Aggregation in an Analytic Task. International Journal of Intelligence and Counter Intelligence. 2014. 27: 752-763.

[8] Seeley T.D, Buhrman S.C 2001 Nest-site selection in honey bees: how well do swarms implement the 'best-of-N' decision rule?. Behav. Ecol. Sociobiol. 49, 416–427

[9] Marshall, James. Bogacz, Rafal. Dornhaus, Anna. Planqué, Robert. Kovacs, Tim. Franks, Nigel. On optimal decision-making in brains and social insect colonies. Soc. Interface 2009.

[10] Seeley, Thomas D., et al. "Stop signals provide cross inhibition in collective decision-making by honeybee swarms." Science 335.6064 (2012): 108-111.

[11] Seeley, Thomas D. Honeybee Democracy. Princeton Univ. Press, 2010.

[12] Seeley, Thomas D., Visscher, P. Kirk. Choosing a home: How the scouts in a honey bee swarm perceive the completion of their group decision making. Behavioural Ecology and Sociobiology 54 (5) 511-520.

[13] Usher, M. McClelland J.L 2001 The time course of perceptual choice: the leaky, competing accumulator model. Psychol. Rev. 108, 550–592

[14] Rosenberg, L.B., "Human Swarms, a real-time method for collective intelligence." Proceedings of the European Conference on Artificial Life 2015, pp. 658-659

[15] Rosenberg, Louis. Artificial Swarm Intelligence vs Human Experts, Neural Networks (IJCNN), 2016 International Joint Conference on. IEEE. J. Clerk Maxwell, A Treatise on Electricity and Magnetism, 3rd ed., vol. 2. Oxford: Clarendon, 1892, pp.68–73.

[16] Ekman, Paul (2014): Spot the Fake Smile: www.ekmaninternational.

978-1-5386-2823-2/17/$31.00 ©2017 IEEE